



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification⁷:

H04Q 11/00, H04L 12/56

A1

(11) International Publication Number:

WO 00/42811

(43) International Publication Date:

20 July 2000 (20.07.00)

(21) International Application Number: PCT/US00/00970

(22) International Filing Date: 14 January 2000 (14.01.00)

(30) Priority Data:

09/232,395

15 January 1999 (15.01.99)

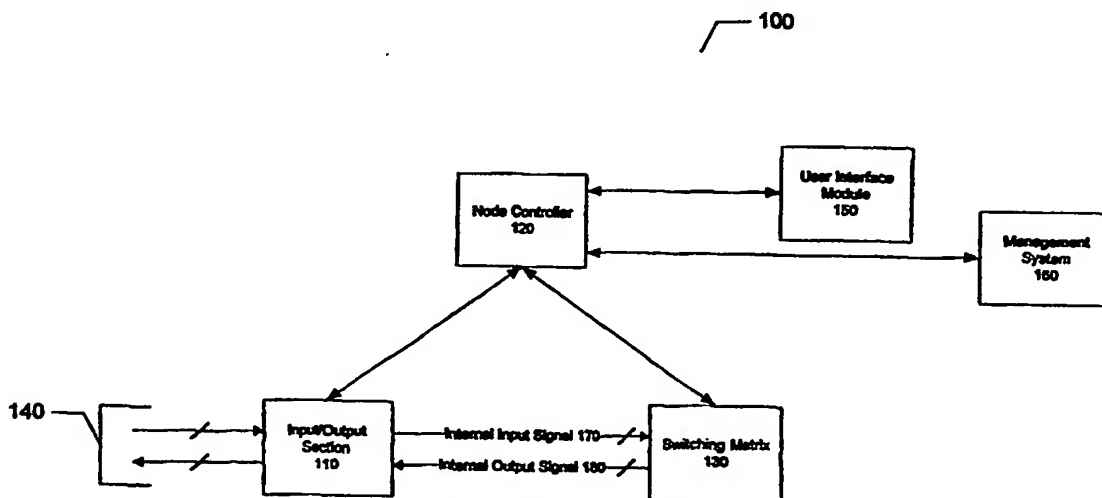
US

(71) Applicant: MONTEREY NETWORKS, INC. [US/US]; 2280
Campbell Creek Boulevard, Richardson, TX 75082 (US).(72) Inventors: ZADIKIAN, Haig, Michael; 919 Parkwood
Court, McKinney, TX 75070 (US). SALEH, Ali, Najib;
31565 Hipshot Drive, Castaic, CA 91384 (US). ADLER,
John, Conlon; 6725 Biltmore Place, Plano, TX 75023
(US). BAGHDASARIAN, Zareh; 238 Valley Creek
Place, Richardson, TX 75080 (US). PARSI, Vahid; 411
Buckingham Road, #534, Richardson, TX 75081 (US).(74) Agents: MACPHERSON, Alan, H. et al.; Skjerven, Morrill,
MacPherson, Franklin & Friel LLP, Suite 700, 25 Metro
Drive, San Jose, CA 95110 (US).(81) Designated States: AU, CA, JP, European patent (AT, BE,
CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC,
NL, PT, SE).

Published

*With international search report.**Before the expiration of the time limit for amending the
claims and to be republished in the event of the receipt of
amendments.*

(54) Title: CONFIGURABLE NETWORK ROUTER



(57) Abstract

A signal router, configured to receive information carried by a first signal and transmit the information on a second signal, is described. The signal router, using routing information it gathers, selects the second signal from a number of signals, with the first and the second signals differing in at least one physical characteristic. The signal router operates in a network of a number of such signal routers, and so each one of the signal routers is coupled to at least one of the other signal routers. The routing information is used to create a circuit from a first one of the signal routers to a second one of the signal routers. Once the circuit is created, the information is routed over the circuit. According to one aspect of the present invention, the signals are optical signals.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

CONFIGURABLE NETWORK ROUTER

TECHNICAL FIELD

5 This invention relates to the field of information networks, and more particularly relates to one or more routers capable of routing information over a network.

BACKGROUND ART

10 Today's networks carry vast amounts of information. High bandwidth applications supported by these networks include streaming video, streaming audio, and large aggregations of voice traffic. In the future, these bandwidth demands are certain to increase. To meet such demands, an increasingly popular alternative is the use of lightwave communications carried over fiber-optic cables. The use of lightwave communications provides several benefits, including high bandwidth, ease of installation, and capacity for future growth.

15 The synchronous optical network (SONET) protocol is among those protocols employing an optical infrastructure. SONET is a physical transmission vehicle capable of transmission speeds in the gigabit range, and is defined by a set of electrical as well as optical standards. SONET's ability to use currently-installed fiber-optic cabling, coupled with the fact that SONET significantly reduces complexity and equipment functionality requirements, gives local and interexchange carriers incentive to employ SONET. Also attractive is the immediate savings in operational cost that this reduction in complexity provides. SONET thus allows the realization of a new generation of high-bandwidth services in a more economical manner than previously existed.

20 SONET networks have traditionally been protected from failures by using topologies that dedicate something on the order of half the network's available bandwidth for protection, such as a ring or mesh topology. Two approaches in common use today are diverse protection and self-healing rings (SHR), both of which offer relatively fast restoration times with relatively simple control logic but do not scale well for large data networks. This is mostly due to their inefficiency in capacity allocation. Their fast restoration time, however, makes most failures transparent to the end-user, which is important in applications such as telephony and other voice communications. The existing schemes rely on 1-plus-1 and 1-for-1 topologies that carry active traffic over two separate fibers (line switched) or signals (path switched), and use a protocol (Automatic Protection Switching or APS), or hardware (diverse protection) to detect, propagate, and restore failures.

30 A SONET network using an SHR topology provides very fast restoration of failed links by using redundant links between the nodes of each ring. Thus, each ring actually consists of two rings, a ring supporting information transfer in a "clockwise" direction and a ring supporting information transfer in a "counter-clockwise" direction. The terms "east" and "west" are also commonly used in this regard. Each direction employs its own set of fiber-optic cables, with traffic between nodes assigned a certain direction (either clockwise or counter

-2-

clockwise). If a cable in one of these sub-rings is damaged, the SONET ring "heals" itself by changing the direction of information flow from the direction taken by the information transferred over the failed link to the sub-ring having information flow in the opposite direction.

The detection of such faults and the restoration of information flow thus occurs very quickly, on the order of 10 ms for detection and 50 ms for restoration for most ring implementations. The short restoration time is critical in supporting applications, such as current telephone networks, that are sensitive to quality of service (QoS) because it prevents old digital terminals and switches from generating red alarms and initiating Carrier Group Alarms (CGA). These alarms are undesirable because such alarms usually result in dropped calls, causing users down time aggravation. Restoration times that exceed 10 seconds can lead to timeouts at higher protocol layers, while those that exceed 1 minute lead to disastrous results for the entire network. However, the price of such quickly restored information flow is the high bandwidth requirements of such systems. By maintaining completely redundant sub-rings, an SHR topology requires 100% excess bandwidth.

An alternative to the ring topology is the mesh topology. The mesh topology is similar to the point-to-point topology used in internetworking. Each node in such a network is connected to one or more other nodes. Thus, each node is connected to the rest of the network by one or more links. In this manner, a path from a first node to a second node uses all or a portion of the capacity of the links between those two nodes.

Networks based on mesh-type restoration are inherently more capacity-efficient than ring-based designs, mainly because each network link can potentially provide protection for fiber cuts on several different links. By sharing the capacity between links, a SONET network using a mesh topology can provide redundancy for failure restoration at less than 100% of the bandwidth capacity originally required. Such networks are even more efficient when traffic transits several links. One study found that for an 11-node, 22-span network, only 51% redundant net capacity was required for 100% restorability, as reported in, "The design and simulation of an intelligent transport network with distributed control," by T. Chujo, H. Komine, K. Miyazaki, T. Ogura, and T. Soejima, presented at the Network Operations Management Symposium, San Diego, February 11-14, 1990, which is incorporated herein by reference, in its entirety and for all purposes. The corresponding ring-based design required five rings and a total DS-3 redundancy of 330%. However, path restoration often consumes several minutes in such a topology. This is much slower than the restoration times exhibited by ring topologies and is so long that connections are often lost during the outage.

Various kinds of networking equipment can be used to support the ring and mesh topologies just described. Options include:

1. Back-to-back wavelength division multiplexers (WDMs) and optical cross-connects (OXCs) for use in mesh topologies.
2. Back-to-back optical add/drop multiplexers (O-ADM) for ring topologies.

-3-

3. Other combinations (e.g., WDM combined with OXC, digital cross-connect systems (DCSs), and other such equipment)

WDMs may be connected in back-to-back configurations to allow the connection of various wavelength routes to one another (also known as "patching" or "nailing up" connections). Provisioning paths in such architectures is done manually using a patch panel. Thus, provisioning is slow and prone to mistakes due to human error and equipment failure. In the event of a failure, restoration is performed manually in such architectures and is again slow and error-prone. Such architectures scale poorly because additional bandwidth is added by either adding to the number of wavelengths supported (requiring the replacement of equipment at nodes, and possibly the replacement of fiber-optic cables as well) or adding new fiber-optic cables and supporting node equipment. Such architectures are also inherently unmanageable, due to the lack of centralized control. And, while the initial capital investment tends to be relatively low (as a result of their simplicity), operating expenses for such architectures tend to be relatively high because of the costs associated with configuration, expansion, and management. Thus, a mesh topology employing back-to-back WDM's will tend to be slow to deploy and difficult to manage due to the need for manually "nailing up" paths and lack of centralization.

Another architectural element that may be used to create a mesh topology is the optical cross-connect (OXC). OXCs allow provisioning using a centralized scheme to accomplish provisioning in a matter of minutes. Restoration in the event of a failure may be performed manually or may be effected using a centralized management system. However, restoration still requires on the order of minutes per wavelength route restored. As with the back-to-back WDM architecture, a mesh topology that employs OXCs scales poorly. This is due in part to the exponential increase in the physical size experienced when expanding the capacity of an OXC with the addition of input and output links. For example, an OXC that supports two links (fiber-optic cables), each having three paths, will need to provide a switching fabric that supports the six possible combinations of connections between the paths carried by the two fiber-optic cables. When this number is increased to four paths per fiber-optic cable, the number of possible connections increases to twenty-four. As still more paths are added to each link and more links are supported, the possible number of connections increases dramatically, increasing the physical size of the affected OXC.

An OXC can be either transparent (i.e., purely optical, in which the signals are never converted to electrical signals) or opaque (i.e., the optical signals are converted into electrical signals and then converted back into optical signals). Transparent optical cross-connects provide little in the way of manageability because the information carried by lightwave is never made accessible to the OXC's operator. In contrast, opaque OXCs can be configured to permit access to the information being switched. However, neither type of OXC maintains information regarding the topology of the network and, in fact, OXCs possess no intrinsic network intelligence. Moreover, OXC technology is expensive, making initial investment quite high, as well as the cost of future expansion.

-4-

Alternatively, a SONET network may be configured in a ring (SHR) topology by using add/drop multiplexers (ADMs). An ADM is a SONET multiplexer that allows DS1 signals to be added into or dropped from an STS-N signal. ADMs have two bidirectional ports, commonly referred to as an east and a west port. Using ADMs, a SONET network in a SHR topology uses a collection of nodes equipped with ADMs in a physical closed loop such that each node is connected to two adjacent nodes with a duplex connection. Any loss of connection due to a single failure of a node or a connection between nodes is automatically restored. The traffic terminated at a failed node, however, is lost. Two types of SHRs are unidirectional (USHR) and bidirectional (BSHR), as defined by the traffic flow in normal conditions. Bidirectional rings have a capacity carrying advantage over unidirectional rings because of the ability to share protection capacity among the links between nodes, as opposed to unidirectional rings, which dedicate capacity all the way around the ring.

Provisioning in such architectures is centralized and can be performed in minutes. While restoration can also be performed quickly (on the order of 50 ms, as previously noted), 100% spare bandwidth is required. Thus, the user must install fiber-optic cabling for two networks, one for normal traffic and one to be used in the event of a failure. Moreover, the cabling for each link should be physically located as far from its corresponding link in order to minimize the possibility that a cause of physical damage will damage both links and cause both directions of a ring to fail. These issues detrimentally affect cost, manageability, and scalability. With regard to expansion, ADMs are stacked in an SHR in order to increase capacity. However, stacked ADMs are blocking. In other words, the switching function may not allow the transfer of data from a port on one stacked ring to a portion on another ring. Thus, an architecture employing ADMs is best suited for small offices or other situations that do not require the relatively large amounts of bandwidth (implying the need for stacked ADMs). As noted, stacked ADMs are also difficult to manage and expensive due to the extra hardware required for 100% spare capacity.

Other combinations can also be employed. For example, WDMs can be combined with OXC's (either transparent or opaque) in order to create a network having a mesh topology. Such an architecture supports the cross-connection of wavelength routes by either manual connection or under centralized control. However, such an architecture is also difficult to expand due to the need to add WDMs/fiber-optic cables and the increase in size of the OXC, and cannot restore failed links quickly enough to avoid dropping or interrupting telecommunications connections.

Another option is the use of a digital cross-connect system (DCS). A DCS is used to terminate digital signals and cross-connect them, integrating multiple functionalities such as signal adding and dropping, cross-connection capabilities, and multiplexing and demultiplexing of signals. DCS based networks enjoy an advantage over networks employing back-to-back WDMs because the use of DCS eliminates the need for additional back-to-back electrical multiplexing, thus reducing the need for labor-intensive jumpers. Operational cost savings are realized by a DCS through electronically controlling cross-connections, test access and loopbacks, and maintenance. Two types of DCSs are wideband DCSs and broadband DCSs. Wideband DCS (W-DCS) terminates full duplex OC-Ns and DS3s, has VT cross-connection capability, and provides DS1 interfaces. A broadband DCS (B-DCS) terminates full-duplex OC-N signals and provides, for example, STS-1 and DS3

interfaces. The B-DCS makes two-way cross-connection at the DS3, STS-1, and concatenated STS-Nc levels. STS-Nc may be used, for example, in broadband services such as high definition television (HDTV), where an STS-3c cross-connection may be used to cross-connect the signal as a single, high-capacity channel.

Various attempts have been made to use DCSs in a mesh configuration to create a fault-tolerant network, but none have been successful in reducing restoration times below a few seconds. Some of these configurations rely on a central database and a central controller (usually an Operations System or OS) to restore failures. Although these schemes often exhibit restoration times exceeding 10 minutes, such restoration times are an improvement over manual restoration, which requires hours, or even days to effect restoration. However, these results are not enough to meet the 50-200 ms restoration time required by existing telecommunication network equipment. Other implementations employ distributed architectures in which control is shared among multiple network nodes. This results in faster restoration times (on the order of about 2-10 seconds), but still does not address the need for restoration times below 200 ms.

DISCLOSURE OF INVENTION

One embodiment of the present invention overcomes conventional limitations by including a network element that supports relatively simple provisioning and relatively fast restoration (on the order of, for example, 50 ms), while providing relatively efficient bandwidth usage (i.e., minimizing excess bandwidth requirements for restoration, on the order of less than 100% redundant capacity and preferably less than 50% redundant capacity). Such a network element is, in one embodiment, based on an architecture that can be easily scaled to accommodate increasing bandwidth requirements.

One embodiment allows a network of nodes, each capable of routing information from one carrier signal to another, to be configured to support the routing of information across the network using those signals to form a circuit. The carrier signals (e.g., optical signals) differ from one another in at least one physical characteristic (e.g., wavelength). The carrier signals, and so the circuit thus selected can be based on routing information obtained from a user, generated by one or more of the nodes, or assembled from other sources. A network element capable of performing such routing functions supports simple provisioning and fast restoration (for example, on the order of 50 ms), while providing efficient bandwidth usage, and can easily scale to accommodate increasing bandwidth requirements.

According to another embodiment of the present invention, a signal router is configured to receive information carried by a first signal and transmit the information on a second signal. The signal router, using routing information, selects the second signal from a number of signals, with the first and the second signals differing in at least one physical characteristic. The signal router operates in a network of a number of such signal routers, and so each one of the signal routers is coupled to at least one of the other signal routers. The routing information is used to create a circuit from a first one of the signal routers to a second one of the signal routers.

Once the circuit is created, the information is routed over the circuit. According to one aspect of the present invention, the signals are optical signals.

According to yet another embodiment of the present invention, a signal router is capable of routing information on different wavelengths. The signal router includes an optical receiver, configured to receive light of a first wavelength. The light of the first wavelength carries the information being routed. The signal router further includes an optical transmitter configured to transmit light of a second wavelength. The optical receiver and the optical transmitter are coupled to a switching matrix. The switching matrix is configured to perform several actions. These actions include selecting the optical transmitter from a number of such optical transmitters, based on routing information obtained by the signal router, and transferring the information from the optical receiver to the selected optical transmitter.

The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting. Other aspects, inventive features, and advantages of the present invention, as defined solely by the claims, will become apparent in the non-limiting detailed description set forth below.

BRIEF DESCRIPTION OF DRAWINGS

The present invention may be better understood, and its numerous objects, features, and advantages made apparent to those skilled in the art by referencing the accompanying drawings.

Fig. 1A is a block diagram of an exemplary router.

Fig. 1B is a block diagram of a network including a number of the routers of Fig. 1A.

Fig. 2 is a block diagram of the signal paths and functional blocks of the router of Fig. 1A.

Fig. 3 is a block diagram of the control paths of the router of Fig. 1A.

Fig. 4 illustrates an exemplary layout of an input/output (I/O) bay.

Fig. 5 illustrates the major components of one of the line cards.

Fig. 6. illustrates an exemplary group matrix.

Fig. 7 illustrates a shelf processor which is responsible for the overall operation, management and control of a shelf.

Fig. 8 illustrates the structure of a multistage matrix.

-7-

Fig. 9 illustrates an example of a physical configuration used for holding one or more matrix stages.

Fig. 10 illustrates one of the switching nodes.

Fig. 11 illustrates a matrix shelf processor.

Fig. 12 illustrates a system controller.

5 Fig. 13 illustrates a route processor .

Fig. 14 illustrates an example of a system switch.

The use of the same reference symbols in different drawings indicates identical items unless otherwise indicated.

MODES FOR CARRYING OUT THE INVENTION

10 The following is intended to provide a detailed description of an example of the invention and should not be taken to be limiting of the invention itself. Rather, any number of variations may fall within the scope of the invention which is defined in the claims following the description.

In addition, the following detailed description has been divided into sections, subsections, and so on, in order to highlight the various subsystems of the invention described herein; however, those skilled in the art will
15 appreciate that such sections are merely for illustrative focus, and that the invention herein disclosed typically draws its support from multiple sections. Consequently, it is to be understood that the division of the detailed description into separate sections is merely done as an aid to understanding and is in no way intended to be limiting.

Fig. 1A illustrates a router 100. Router 100 includes an input/output section 110, a node controller 120,
20 and a switching matrix 130. Node controller 120 contains, for example, real time software and intelligent routing protocols (not shown). Router wavelength 100 supports interfaces including, but not limited to, optical signal interfaces (e.g., SONET), a user interface module 150, and a management system 160. Internal input signals 170 and internal output signals 180 may be electrical or optical in nature. Fig. 1B illustrates a network 190 that
25 includes a number of nodes, network nodes 195(1)-(N). One or more of network nodes 195(1)-(N) can be a router such as router 100. Network 190 can thus support the automatic provisioning, testing, restoration, and termination of virtual paths (exemplified by a virtual path 191) over a physical path (exemplified by a physical path 192) from one of network nodes 195(1)-(N) to another of network nodes 195(1)-(N).

Among other benefits, router 100 solves three growth-related problems often encountered in today's information networks, and particularly in SONET networks:

1. Port Capacity growth: Router 100 includes, for example, a scaleable architecture which can provide, for example, (i) 250 or more nodes/network and (ii) at least 4096 ports/nodes, at a relatively low cost and high density.
2. Bandwidth management: The distributed management architecture of one embodiment of exemplary router 100 allows some or all nodes in the network to be managed from a single workstation. Provisioning a new connection is easily accomplished. Provisioning may be effected, for example, by selecting the source and destination nodes and specifying the required bandwidth and desired quality of service (QoS). An QoS-based shortest-path first (SPF) path selection method is invoked to calculate the best route for the new connection. The QoS-based technique can take into consideration parameters such as existing trunk allocations, network status, the priority and desired quality of the new connection, and other such criteria. This can be accomplished, for example, by sending one or more configuration requests to, and awaiting acknowledgment replies from, the nodes along the new connection's path.
3. Efficient and fast restoration: An exemplary network of two or more routers 100 as illustrated in Fig. 1B preferably uses a mesh topology. Through the use of the routers 100, the network can be more efficient than existing ring topologies, especially when connections span multiple rings. This is possible because a single spare connection in a mesh network can provide protection for several different possible span cusps. In a ring-based network, however, spares can typically only protect against failures on their own ring. Furthermore, when connections span multiple rings, "dedicated" spare bandwidth must typically be allocated on every ring along the path. No other connections can share this spare bandwidth. Recent studies have found that mesh restoration typically requires only about 51% redundancy to yield 100% restorability, while a corresponding ring-based design typically require 330% redundancy for 100% restorability. One embodiment of router 100 supports the restoration of a majority of network failures within less than 50 ms, thus eliminating an advantage that rings generally have over mesh topologies: fast restoration time. A protocol, such as that according to the co-pending application entitled "A METHOD FOR ROUTING INFORMATION OVER A NETWORK" (as previously referenced), can be run on such a router and encompasses all aspects of the restoration process: alarm gathering, path implementation (including alternate path discovery), and path assurance. In cases where there is insufficient bandwidth to satisfy all failed connections, the protocol, in one embodiment, can use a quality of service (QoS) metric to prioritize the restoration sequence. In such embodiment, connections with the highest QoS are restored first, followed, in a descending order, by those with a lower QoS, until either all connections have been restored or all available bandwidth has been used.

Router 100 is a multi-rack, fully redundant router that, in one embodiment, supports at least 256, 1+1 I/O ports, and provides 1-plus-1 protection by using multiple copies (e.g., two or more) of group and main matrices operating in 1+1 mode. Failures within one copy of a given matrix do not require a complete switchover to the backup copy. Only the affected paths through the matrix are switched to the backup copy. This greatly improves switching speed and minimizes the impact of such redundancy on other connections. Preferably, the

group matrix is a 2:1 reduction stage that selects output signals from one of two line cards or I/O modules and connects the selected output signals to the main matrix, thus preventing non-working antecedent from consuming any ports on the main matrix.

5 In one embodiment, there are at least three types of processors in a router 100. The lowest level, level-3, resides on the line card, also referenced to herein as an I/O module, and is responsible for all real time aspects of the processing of the physical protocol (e.g., SONET). In a SONET implementation, every level-3 processor is responsible for a single optical signal (e.g., an OC-48 signal) and, via a protocol processor, performs all required SONET/SDH section and line termination functions. The fast response time required from the level-3 processor makes a firmware implementation preferable. The firmware, which may be written in the "C" or "C++"
10 programming languages, assembler, or other programming language, is preferably optimized for low latency and resource efficiency. Higher-level processing is implemented on a separate module, the shelf processor module, which is shared by several line cards.

The second level of processors, level-2, reside on a shelf and main matrix processor modules. The software on the shelf processor module is responsible for managing and controlling line cards. Only half the line
15 cards supported are active at any one time in order to support 1+1 protection. A level-2 processor deals with tasks that require a reasonable response time (for example, on the order of milliseconds), but have no direct impact on the data path. In other words, missed events, such as hardware interrupts, do not result in bit errors. Some of the functions handled by the shelf processor include the periodic collection of maintenance data from the line cards, receiving and processing periodic keep-alive messages from those cards, shelf startup and
20 configuration, proxy management, and other related functions.

The third processor level, level-1, resides on a system processor module and provides system-wide management and control services. In one embodiment, there are preferably two fully synchronous copies of the level-1 processor in the system, both of which are simultaneously active and, through a dedicated and redundant high-speed link, keep their run-time and stored databases fully synchronized. One of the two processors is
25 designated the master and is responsible for all level-1 processing. An update message is sent to the second processor whenever a change is made to the database and before that change is effected. A periodic keep-alive mechanism allows either copy of the system controller to detect failures on the other copy.

Router 100 provides yet another type of processor, referred to herein as a route processor. Such a processor is dedicated to the path/route discovery and restoration functions. The route processor is responsible
30 for receiving failure indications from the line cards, calculating a new route for failed connections, and sending reconfiguration requests to all affected nodes, including its own.

Hardware Architecture

In one embodiment, router 100 is a multi-rack communications system capable of terminating at least 8192 signals and cross-connecting at least 4096 OC-48 signals. Such a router can be used, for example, as

-10-

SONET/SDH line terminating equipment (LTE) capable of terminating the Section and Line overheads of received OC-48 signals, and cross-connects those signals according to provisioned input-output mappings. Some of the terminated signals can optionally be protected using any of the common protection schemes (1+1, 1:1, and 1:N).

- 5 Overhead processing and generation is performed on the line card by a protocol processor. This protocol processor handles all aspects of the SONET protocol, including framing, insertion and extraction of embedded data channels, error checking, AIS detection, pointer processing, clock recovery, multiplexing/duplexing, and similar duties.

Signal Path

- 10 Fig. 2 is a block diagram of signal paths 200 within router 100. The primary signal paths in router 100 include one or more groups exemplified by groups 210(1)-(N), group matrices 212(1)-(N), and a main matrix 214. As depicted in Fig. 1A, groups 210(1)-(N), and group matrices 212(1)-(N) are shown as having receive and transmit sections. Groups 210(1)-(N) each include line cards 220(1,1)-(1,N), through line cards 220(N,1)-(N,N). Signals from line cards 220(1,1)-(N,N) are sent to the corresponding group matrix. In one embodiment, two sets
15 of the group matrix cards, group matrices 212(1)-(N) and group matrices 216(1)-(N) are employed. Main matrix 214 is also mirrored in one embodiment by a redundant copy, a backup main matrix 218, which together form switching matrix 130. As shown in Fig. 2, the redundancy for group matrices 212(1)-(N) (i.e. group matrices 216(1)-(N)), is also provided on the transmit side.

- NOTE: The variable identifier "N" is used in several instances in Fig. 2 (and subsequent use of other
20 variables, such as "m," "x," "k," and others) to more simply designate the final element (e.g., group matrix 212(N), line card 220(N,N), and so on) of a series of related or similar elements (e.g., group matrices 212(1)-(N), line cards 220(1,1)-(N,N), and so on). The repeated use of such variable identifiers is not meant to imply a correlation between the sizes of such series of elements. The use of such variable identifiers does not require that each series of elements has the same number of elements as another series delimited by the same variable
25 identifier. Rather, in each instance of use, the variable identified by "N" (or "m," "x," "k," and others) may hold the same or a different value than other instances of the same variable identifier. For example, group matrix 212(N) may be the tenth group matrix in a series of group matrices, whereas line card 220(N,N) may be the forty-eighth line card in a series of line cards.

- Using signal paths 200 as an example, data enters the system at one of line cards 220(1,1)-(N,N). It is at
30 this point, in a SONET-based system, the Section and Line overheads are processed and stripped off by a protocol processor (not shown). The extracted SONET/SDH payload envelope is then synchronized with the system clock and sent to two different copies of a local matrix, depicted as group matrices 212(1)-(N) and 216(1)-(N) in Fig. 1A. In one embodiment, group matrices 212(1)-(N) and 216(1)-(N) are used mainly as 2:1 reduction stages that select one of two optical signals and pass the selected optical signal to switching matrix 130. This allows the

-11-

implementation of a variety of protection schemes (including 1:N, or 0:1) without having to use any additional ports on main matrix 214. All protect signals are terminated at group matrices 212(1)-(N) and 216(1)-(N). In order to maximize bandwidth, it is preferable that only active signals be passed through to switching matrix 130.

In one embodiment, switching matrix 130 is an errorless, rearrangeably non-blocking switching network.

5 In one embodiment, switching matrix 130 is a 256x256 switching network that consists of three columns and 16 rows of 16x16 switching elements that allow any of their inputs to be connected to any of their outputs. Also, preferably a single copy of the matrix is housed in a single rack that contains three shelves, one for each column (or stage) of the matrix. Each shelf contains cards housing the 16 switching elements in each stage. The switching element itself may include, for example, a 16x16 crosspoint switch, with optical transceivers, and a

10 microcontroller for controlling the crosspoint switch and providing operational feedback to the level-2 processor. Communications between the two processors may be carried, for example, over an Ethernet connection. The level-2 processor in turn communicates with the level-1 and route processors using, for example, a redundant Ethernet connection.

The switching elements in each matrix copy of the exemplary embodiment may be connected using fiber-

15 optic cables, for example. While copper cabling may also be employed, such an option may not offer the speed and number of connections provided by an optical arrangement. After passing through the stages of switching matrix 130, an optical signal may be routed to an I/O shelf that (optionally) splits it into two signals. One of the signals is sent to an active line card, while the other, when available, is sent to a backup card.

Line cards 220(1,1)-(N,N) receive optical signals from group matrices 212(1)-(N) and 216(1)-(N) which

20 are in turn connected to two separate copies of the main matrix. Line cards 220(1,1)-(N,N) monitor both signals for errors and, after a user-defined integration period, switch to the backup signal if that signal exhibits better bit error rate (BER) performance than the prior active signal. This scheme, referred to herein as 1-plus-1, allows line cards 220(1,1)-(N,N) to select between the two copies of the group matrix without any level-1 or level-2 CPU intervention. This helps to ensure that such a switch can be made in 50 ms or less (per Bellcore's

25 recommendations in GR-253 (GR-253: *Synchronous Optical Network (SONET) Transport Systems*, Common Generic Criteria, Issue 2 [Bellcore, Dec. 1995], included herein by reference, in its entirety and for all purposes)). The selected signal is then processed by the transmit section of the protocol processor, which inserts all required transport overhead bytes into the outgoing stream.

Regarding the signals described herein, both above and subsequently, those skilled in the art will

30 recognize that a signal may be directly transmitted from a first logic block to a second logic block, or a signal may be modified (e.g., amplified, attenuated, delayed, latched, buffered, inverted, filtered or otherwise converted, etc.) between the logic blocks. Although the signals of the embodiments described herein are characterized as transmitted from one block to the next, other embodiments may include modified signals in place of such directly transmitted signals with the informational and/or functional aspect of the signal being transmitted between blocks.

35 To some extent, a signal input at a second logic block may be conceptualized as a second signal derived from a

-12-

first signal output from a first logic block due to physical limitations of the circuitry involved (e.g., there will inevitably be some attenuation and delay). Therefore, as used herein, a second signal derived from a first signal includes the first signal or any modifications to the first signal, whether due to circuit limitations or due to passage through other circuit elements which do not substantively change the informational and/or final functional aspect of the first signal.

Control Path

Fig. 3 illustrates a control path 300 of a router, such as router 100. Control path 300 includes all non-payload-related flows within the system and the hardware and software necessary to the control of the signal paths illustrated in Fig. 2. All major control flows are carried over an internal local area network (LAN), which is, for example, a collection of switched Ethernet segments. The structure of the internal LAN is hierarchical and can be created using a mixture of 10Mbps and 100Mbps Ethernet segments, for example. Higher-speed segments (e.g., gigabit Ethernet) can be used as well.

Groups

At the bottom of the hierarchy is what is referred to herein as a group matrix, or a Group Ethernet Repeater in a system using Ethernet communications, and depicted in Fig. 3 as group matrices 212(1)-(N) and 216(1)-(N). Each one of group matrices 212(1)-(N) and 216(1)-(N), also referred to herein as a hub, a repeater, or concentrator, is a physical layer device and preferably supports a star network topology, such as the IEEE 802.3 10BASE-T networking standard. The redundant connections from line cards 220(1,1)-(N,N) in each of groups 310(1)-(N) are connected to two repeaters that reside on two separate copies of the group matrix module. Preferably, each one of line cards 220(1,1)-(N,N) supports two network ports (e.g., 10BASE-T Ethernet ports). The two sets of four signals from each port pass through a relay that selects one of them for connection to the LAN for purposes of redundancy. Groups 310(1)-(N) represent the first layer of the control bus hierarchy. Group matrices 212(1)-(N) and 216(1)-(N) are each controlled by a shelf processor (not shown, for the sake of clarity) and communicate with one of the shelf switches described below via LAN connections.

Shelf Ethernet Switch

Fig. 3 also illustrates certain features of router 100 pertaining to the relationship between shelf switches 320(1)-(N) and 321(1)-(N), and groups 310(1)-(N). Groups 310(1)-(N) are again shown, with regard to the control functions thereof. In this depiction of groups 310(1)-(N), line cards 220(1,1)-(N,N) are shown as being attached to networking devices, indicated here as group matrices. Group matrices 212(1)-(N) and 216(1)-(N) may be, for example, multi-port Ethernet hubs running at 10Mbps. Each of line cards 220(1,1)-(N,N) feed signals into

-13-

two of group matrices 212(1)-(N) and 216(1)-(N). For example, line card 220(1,1) feeds received information to group matrix 212(1) and group matrix 216(1). Group matrices 212(1)-(N) and 216(1)-(N) each feed a signal into shelf switches 320(1)-(N) and 321(1)-(N) of Fig. 2. Shelf switches 320(1)-(N) and 321(1)-(N) are each controlled by a shelf processor (not shown, for the sake of clarity) and communicate with one of the system switches (not shown, for the sake of clarity).

Shelf switches 320(1)-(N) and 321(1)-(N) are the next higher level of the control hierarchy in router 100, and are located on the shelf processor module (exemplified by line racks (330(1)-(N))). Each copy of shelf switches 320(1)-(N) and 321(1)-(N) interconnects six connections from the three groups in each shelf, another connection from the shelf processor, and one connection from system switch 340 (and 341). Shelf switches 320(1)-(N) and 321(1)-(N) can be implemented, for example, using an 8-port Ethernet configured to handle 10Mbps Ethernet traffic and a single-port, dual-rate switch (e.g., 10Mbps/100Mbps Ethernet).

System Switch

The next level of the hierarchy is the system switch (in routers using Ethernet-based inter-processor communications, this is referred to as the system Ethernet switch), of which there are two copies in each router. These are shown as system switches 340 and 341 in Fig. 3. This fully redundant scheme prevents failures on one shelf switch from taking down the entire control bus. In one embodiment, a system switch manages connections from the following sources:

1. High-speed connection(s) from shelf switches 320(1)-(N) and 321(1)-(N);
2. High-speed connection(s) to higher-level processors (e.g., redundant level-1 processors 350 and 351, and redundant route processors 360 and 361); and
3. High-speed connection(s) to matrix shelf processors 370(1)-(N) and 371(1)-(N) which, in turn, control matrix cards 380(1,1)-(1,N)), located in main matrix racks 390(1)-(N).

It will be noted that main matrix 214 includes matrix cards 380(1,1)-(1,N), and that, more generally, main matrices 214 and 218 are included matrix racks 390(1)-(N)

System switches 340 and 341 are located in a management bay. As noted, the fully redundant switches manage connections from various router elements, such as I/O and matrix bays, level-1 processors, and route processors. Each of level-1 processors 350 and 351 and route processors 360 and 361 is preferably connected to system switches 340 and 341 using 100Mbps Ethernet connections in a configuration that creates an expandable, efficient, and fully redundant control bus. If more inter-processor communication bandwidth is required, then the connection is preferably a higher speed connection, such as that provided by a gigabit Ethernet or fiber-channel connection.

Physical configurations and modules

I/O Bay

Fig. 4 illustrates an exemplary layout of an input/output (I/O) bay 400. The I/O bay shelf can support, for example, a total of 16 slots. Slots may be logically divided into functional groups. In such an embodiment, four such functional groups are defined with three of the groups occupying five slots each. In that embodiment, the other group, which occupies a single slot can be configured to house the shelf processor. Thus, in the embodiment shown in Fig. 4, I/O bay 400 contains line cards (LC) 410(1)-(N), group matrices (GM) 420(1)-(N), which are controlled by shelf processors (SP) 430(1)-(N), which are exemplary of line cards 220 (1,1)-(N-N), group matrices 212 (1)-(N) and 216 (1)(N), and shelf processors 320 (1)-(N) and 321(1)-(N), and shelf switches 440(1)-(N). It will be noted that the various line cards, group matrices, and shelf processors correspond to similar elements from previous figures.

Groups and Magazines

A group is made up of line cards occupying a number of slots on a shelf. A slot is also referred to herein as a magazine. In one implementation, the group is 20 line cards that occupy five slots. Four of the slots hold, for example, 16 line cards at 4 per slot. The same slot can be used with a wide variety of I/O modules and in various configurations. One example of this flexibility, in a SONET configuration, is the ability to house an OC-192 I/O modules in the same space occupied by four OC-48 line cards. In fact, the slots in each group are not required to be of the same type or structure. This architecture provides flexibility to allow any combination of line cards to be installed in each slot.

The fifth slot in the aforementioned embodiment can be configured to accept line cards containing an optical switching matrix and a hub (e.g., an Ethernet hub). Preferably, two group matrix cards are employed, each containing a 2:1 optical reduction stage that "selects" working channels before the signals leave the shelf. In a 1+1 protection scheme, the two inputs to the line cards are classified as active and protect channels. The working channel is one of the active and protect channels that is selected based on bit error rate or other criteria, and so implements a redundancy scheme. This prevents the standby line cards from using any bandwidth on switching matrix 130.

Backplane

The following describes one embodiment of a backplane and some of the interface signals on that backplane. The backplane in the I/O bay shelf carries a variety of signals between line cards and other modules in the shelf. Each I/O shelf module is configured to allow an automatic, errorless switch from one power bus to the other. Backplane signals that are common to all modules in the I/O shelf includes power, ground, and signal ground.

-15-

Shelf processor module backplane signals include reset signals, clock signals, hardware detect signals (e.g., card detect, copy present, and the like), slot ID signals, and slot communication signals (both low and high speed). I/O module (line card) backplane signals include reset signals, clock signals, communication signals, hardware detect signals, and slot ID signals. Group matrix module backplane signals include reset, clock signals, communication signals (both low and high speed), detection and hardware detect signals, and slot ID signals.

System Modules

Line Card

Fig. 5 illustrates the major components of one of line cards 220(1,1)-(N,N), exemplified in Fig. 5 by a line card 500. A line card, also referred to herein as an I/O module, integrates all the necessary hardware and software functions to properly terminate the physical layer. In a SONET implementation, a line card terminates the transport overhead (Section + Line) of a full duplex OC-48 signal. Other components on this card provide a redundant optical connection to the switch matrix, and a communication channel to other modules in the system.

Line card 500 receives optical signals from other network elements via a line-side optical receiver 505 and from the local router's system via a system-side optical receiver 506. Each of these receivers implements an optical-to-electrical (O/E) conversion function. Line card 500 transmits optical signals to other network elements using a line-side optical transmitter 510 and to the group matrices using a system-side optical transmitter 511. Each of these transmitters implements an electrical-to-optical (E/O) conversion function. It will be noted that line-side refers to the side of the line card coupled to other network elements and system-side refers to the side of the line card coupled to the group matrices.

Line-side optical receiver 505 is coupled to a protocol processor 520 which performs clock recovery multiplexing, demultiplexing, and SONET STE/LTE processing in both directions. Similarly, system-side optical receiver 506 is also coupled to protocol processor 520 to allow protocol processor 520 to receive optical signals. The processed electrical signals from protocol processor 520 are coupled to the transmitters 510 and 511. The clock recovery functions are combined with demultiplexers and multiplexers to support reception and transmission of the optical data, respectively. The multiplexers serialize output data generated in protocol processor 520 by performing parallel-to-serial conversion on the parallel data. In contrast, de-multiplexers are used in protocol processor 520 to perform serial-to-parallel conversion on received data.

In order to add protection channels, line-side optical transmitter 510 is also coupled to a 1:2 broadcast unit 535. To receive such optical signals, optical receiver 506 is also coupled to a 2:1 selector 536 in order to select the working channel before the optical signals leave the shelf and thus prevent the standby channel (also referred to herein as the protect channel) from using any bandwidth on switching matrix 130.

Protocol processor 520 is coupled to a bus 545. Protocol processor 520 interfaces line card 500 to two copies of the matrix in a 1+1 physical protocol. In a SONET implementation, protocol processor 520 provides

-16-

both STE/LTE processing according to published industry standards. Also coupled to bus 545 are a memory 560 and a CPU 570. Memory 560 should be fast enough for efficient operation of CPU 570.

CPU 570 communicates with other of line cards 220(1,1)-(N,N) over a control bus (not shown) using a transceiver 580 that is coupled to CPU 570. Transceiver 580, is coupled to a transformer 585 which is coupled to a switch 590. Switch 590 is coupled to the control bus. Switch 590 implements a 1:1 protection scheme for transceiver 580 and couples CPU 570 to two independent ports on the backplane (not shown). Each of the two ports connects to one copy of the hub of the group matrix. This allows the software on the line card to switch to the backup link when it detects failures on the active link.

Preferably, CPU 570 includes numerous integrated peripherals including embedded SCC channels (e.g. M-band communications) and an Ethernet controller (for example, to support communications with other system modules). In one embodiment, CPU 570 provides an onboard communications processor module (not shown) that handles time-critical aspects of the protocols supported.

Group Matrix Module

The group matrix module includes two independent blocks: a group matrix and a hub (also referred to herein as a repeater).

Group matrix

Fig. 6. illustrates an exemplary group matrix 600, which is exemplary of group matrices 212(1)-(N) and group matrices 216(1)-(N). In the embodiment shown in Fig. 6, group matrix 600 includes a series of 2:1 path selectors (exemplified by selectors 610(1)-(N)), broadcast units 620(1)-(N), and a microcontroller 630 controlling these. Selectors 610(1)-(N) select one of two full-duplex optical signals and couple the selected signal to switching matrix 130. Selectors 610(1)-(N) and broadcast units 620(1)-(N) are grouped into pairs to form I/O channels 645(1)-(N). Microcontroller 630 communicates with other elements of router 100 via redundant transceivers (exemplified by transceivers 635 and 640). For example, microcontroller 630 can control selectors 610(1)-(N) and broadcast units 620(1)-(N) through commands received from the group processor.

Hub

One or more hubs are also provided to support communication between the group matrices and system switches in router 100. In an Ethernet communications environment, the hub's functions are carried out primarily by repeater interface controllers (RICs). Each RIC integrates the functions of a repeater, clock and data recovery unit (CDR), Manchester encoder/decoder, and transceiver. Each RIC has a set of registers that convey status information and allow a number of configuration options to be specified by the user using, for example, a microcontroller.

Shelf Processor Module

-17-

The shelf processor module provides, among other elements, a shelf processor and switch that interconnect the LAN segments from the groups and the shelf processor to a port on the shelf switch (Ethernet switch 730).

Shelf Processor

5 Fig. 7 illustrates a shelf processor 700 which is responsible for the overall operation, management, and control of the shelf. A shelf CPU 705 controls the functions of shelf processor 700. Shelf CPU 705 is connected to a debug port 707 via a debug port transceiver 710. Debug port 707 may be a device capable of coupling shelf CPU 705 to a personal computer or dumb terminal. Debug port 707 allows a user to access shelf processor module 700 to determine the cause of any errors therein. Transceivers 711 and 712 each connect an SCC channel
10 of shelf CPU 705 to the other shelf processor. The resulting link, which can use high-speed asynchronous framing, serves as an inter-processor communications interface.

Shelf CPU 705 is also connected to a timer 715, which preferably contains the following three functional blocks:

- 15 1. Power-fail-reset: Monitors the supply voltage and restarts the processor when power failures occur and generates a power-up reset pulse.
2. External reset: Provides a push-button interface that debounces the input signal and provides a reset pulse.
3. Timer: An internal timer that generates a reset pulse if the strobe input signal is not toggled prior to timeout (e.g., 150 ms, 600 ms, or 1.2 seconds).

20 Shelf CPU 705 also accesses a memory 721 and a reset latch 722 over a CPU bus 725. Reset latch 722 supports reset of one or more line cards (not shown). Shelf CPU 705 is also coupled to an Ethernet switch 730. The network switch interconnects the lower speed inter-processor communication network segments in each shelf. In one embodiment, the network switch provides support for 10Mbps and 100Mbps segments. In one embodiment, an integrated bus master and slave interface allow multiple devices to be interconnected.

25 Ethernet switch 730 is coupled to a transceiver 735 which, via a select 740, allows Ethernet switch 730 to connect to two separate Ethernet segments. Select 740 implements a 1:1 protection scheme that allows shelf processor 700 to recover from failures on the active segment by simply switching to the other segment. Ethernet switch 730 is also coupled to one or more group transceivers (exemplified by group transceivers 750, 751, 752, and 753). Group transceivers 750, 751, 752, and 753 connect ports on Ethernet switch 730 to the groups.

30 System Switch

One embodiment of a system switch (or system Ethernet switch, in routers that communicate using Ethernet) capable of interconnecting at least 13 network segments in a switched configuration. In an Ethernet-

-18-

based system, the system switch supports both 10 Mbps and 100 Mbps connections. The segments come from the shelf switching in the I/O shelf and the matrix switches, among others, and the system switch allows these elements to communicate.

Main Matrix Bay

5 A switching matrix in router 100 is based on a rearrangeable non-blocking network. A switching matrix, as described herein consists of switch nodes arranged in a staged array. For a 256x256 switching matrix, for example, switch matrix 130 consists of 48 nodes arranged in an array of 16 rows by 3 columns, with each column containing one stage of the switch matrix. All 48 nodes in the switch matrix are substantially similar and consist of a 16 x 16 crossbar device that allows any of its 16 inputs to be connected to any of its 16 outputs, regardless of
10 the current state of the crossbar.

Main Matrix

Fig. 8 illustrates switching matrix 130 configured in the manner of the switch matrix just described. In one embodiment, switching matrix 130 employs a 256x256 matrix, an array of switching nodes 800(1,1)-(16,3), each of which is a 16x16 crossbar switch that allows any of the 16 input signals to be connected to any of its 16
15 outputs, regardless of the current state of the crossbar. In one environment, each of the interconnections between switching nodes 800(1,1)-(16,3) represent dual gigabit interconnections. As noted, the embodiment illustrated in Fig. 8 supports the switching of up to 256 inputs, shown as inputs 820(1)-(256). Inputs 820(1)-(256) are switched to one of outputs 830(1)-(256). Physically, each of the 48 switching nodes of this embodiment occupies a single slot in the matrix rack. The rack itself, which is shown in Fig. 9, is made up of three shelves (one per matrix
20 column) that house the switch node cards (there are 16 such cards in every shelf) and six-shelf-processor cards (two per shelf).

Matrix Rack

Fig. 9, as noted, illustrates an example of a physical configuration used for holding one or more matrices, and referred to herein as a matrix rack 900. In one embodiment, matrix rack 900 is configured to hold the 48
25 switching nodes (i.e., switching nodes 800(1,1)-(16,3)) in a physical configuration as illustrated in Fig. 10. Switching nodes 800(1,1)-(16,3) are configured as indicated in Fig. 10. Each of switching nodes 800(1,1)-(16,3) supports 16 input signals and 16 output signals, and thus provides switching matrix 130 with 256 input signals and 256 output signals as shown in Fig. 8, 830(1)-(256). Also shown in Fig. 9 are matrix shelf processors 910(1)-(16). Matrix shelf processors 910(1)-(6) are configured in redundant pairs to provide fault-tolerant control of
30 switch nodes 800(1,1)-(16,3). Thus, matrix shelf processors 910(1) and 910(2) control the first "column" (i.e., switching nodes 800(1,1)-(16,1)), matrix shelf processors 910(3) and 910(4) control "column 2" (i.e., switching nodes 800(1,2)-(16,2)), and matrix shelf processors 910(9) and 910(6) control "column 3" (i.e., switching nodes 800(1,3)-(16,3)).

-19-

The cross-connect information, i.e. input-to-output mapping, is written into the crosspoint switch by a local microcontroller which receives it from the local shelf processor over a high-speed connection. The three shelf processors in each rack receive such information from the node controller, which resides in a different rack. This hierarchy can be extended indefinitely. The crosspoint switch receives a high speed serial data from the optical receivers that perform optical-to-electrical conversion on the received optical signals. Data from the crosspoint switch is re-timed to synchronize the data _____ with the system clock of router 100, using a clock and data recovery (CDR) unit, before being converted back into an optical signal that connects to the next stage of the matrix over fiber-optic cables.

Switch Node Module

Fig. 10 illustrates one of switching nodes 800(1,1)-(16,3) as a switching node 1000. Switching node 1000, in one embodiment, is a complete, strictly non-blocking, 16x16 OC-48 multi-stage crossbar matrix which allows any of its inputs to be connected to any of its outputs regardless of the current state of the matrix. A crosspoint switch 1005 is controlled by a local microcontroller (a microcontroller 1010) that also manages the optical transceivers, CDRs, and onboard SONET device. Switch node 1000 configuration is downloaded from microcontroller 1005 over a low-speed bus.

The block diagram of switch node 1000 in Fig. 10 illustrates the main elements of a switch node using a SONET-based implementation. The core of the switch node 1000 is crosspoint switch 1005, which is a 16x16 crossbar switch when implementing a 256x256 matrix. Crosspoint switch 1005 is preferably a 2.5Gbps 16x16 differential crosspoint switch with full broadcast capability. Any of its input signals can be connected to any, or all, of its output signals. The device is configured through a low-speed port that, through a two-step/two-stage process, allows changes to be made to switch configuration without disturbing its operation.

Assuming 16 input signals (indicated in Fig. 10 as inputs 1015(1)-(16)), crossbar switch 1010 is configured to receive optical input signals from optical receivers 1020(1)-(16) at switch input signals 1021(1)-(16). Crossbar switch 1010 also provides switch outputs 1022(1)-(16), which serve as the source of optical output signals for switch node 1000. Microcontroller 1010 is also responsible for detecting and reporting loss-of-signal (LOS) and out-of-lock (OOL) conditions from the optical receivers and CDRs, respectively. Microcontroller 1010 communicates with the shelf processor via transceivers 1060 and 1065 over a bus that carries asynchronous data over the backplane (not shown).

Incoming signals are routed to one of switch outputs 1022(1)-(16) by crosspoint switch 1005 under the control of microcontroller 1010. Switch outputs 1022(1)-(16) are coupled to CDRs 1070(1)-(16), which in turn drive optical transmitters 1080(1)-(16). Output signals from optical transmitters 1080(1)-(16) appear at outputs 1090(1)-(16) as optical signals.

Matrix Shelf Processor Module

Matrix shelf processor 1100 module provides local control and management for one of the main-matrix shelves. The matrix shelf processor 1100 communicates with the level-1 and route processors over a low speed network connection and with the matrix node cards over a multi-drop, low-speed bus.

Fig. 11 illustrates a matrix shelf processor 1100, which is illustrative of matrix shelf processors 910(1)-(6) of Fig. 9 and shelf processor 700 of Fig. 7. Matrix shelf processor 1100 provides local control and management for one of the shelves of a main matrix such as switching matrix 130 (Fig. 1). The core of matrix shelf processor 1100 is a matrix shelf processor CPU 1110. Matrix shelf processor CPU 1110 communicates with one or more level-1 processors (not shown) and route processors (not shown) via a transceiver 1120 (preferably a 10BASE-T transceiver). Matrix shelf processor CPU 1110 communicates with the system switches (i.e., system switches 340 and 341) via a transceiver 1140. To support these functions, matrix shelf processor CPU 1110 is coupled via a processor bus 1170 to memory 1160 which provides storage for various software modules run by matrix shelf processor CPU 1110.

Management Bay

The management bay can house, for example, the following modules:

1. Level-1 processors, or system controllers, and their associated storage devices;
2. Route processors;
3. Optional group and WAN cards that provide high-speed (e.g., greater than T1) X.25 links to one or more operations systems (OS's);
4. System Ethernet switches; and
5. Synchronization modules.

All of the above modules are fully redundant and communicate with the rest of router 100 over redundant control buses. The placement of individual modules within the rack is not addressed in this document, since there are no architectural preferences, or restrictions, on such choices.

Level-1 Processor/System Controller

Fig. 12 illustrates a system controller 1200 (also referred to herein as a level-1 processor). The core of the system controller 1200 is a processor 1210, which also communicates with the system switches (i.e. system switches 340 and 341). Programs run on processor 1210 are stored in memory 1220 coupled thereto. Processor 1210 is also coupled to an all-purpose bus (APB) 1230, which in turn drives several bus and communications controllers. Among the controllers interfaced to APB 1230 is a bus bridge 1240, a peripheral interface 1250, and

-21-

an I/O interface 1260. I/O interface 1260 may provide functionality such as 10Mbps/100Mbps Ethernet communications. I/O interface 1260 also supports peripherals such as keyboards, mice, floppy drives, parallel ports, serial ports, and the like. Bus bridge 1240 allows communications between processor 1210 and other devices. Peripheral interface 1250 allows communications with peripherals such as hard disks. The level 1 processor performs various functions, such as communicating with the route processor(s) to determine how the matrix should be configured, managing the router's resources, and similar duties.

APB 1230 may also be connected to a dual-channel serial communication controller (SCC), which is used to communicate with one or more remote Operations Systems (OS) using, for example, the X.25 protocol. For more OS links and higher link speeds, the user can optionally install one or more WAN Interface Modules in the management bay. Such modules, which preferably handle all real-time aspects of the OS link, including layer-2 of the OSI stack, communicate with the level-1 processor.

Route Processor Module

Fig. 13 illustrates a route processor 1300. Route processor 1300 is a high-speed processor subsystem with relatively limited I/O capabilities. Route processor 1300 functions to receive link-failure indications from the line cards (not shown), computes an alternate route for failed connections using a restoration protocol such as that described in the co-pending application entitled "A METHOD FOR ROUTING INFORMATION OVER A NETWORK" and previously included by reference herein, and then sends one or more configuration requests to all affected nodes to achieve this new routing. Route processor 1300 is able to communicate directly with all system modules, including the line cards (not shown) and the matrix shelf processors (not shown) via a redundant high speed network connection to the system switch. In systems using Ethernet as the communication mechanism, route processor 1300 communicates with these elements via a redundant 100Mbps connection to the system Ethernet switch. The core of route processor 1300 is a processor 1310 which runs software stored in memory 1330 via a CPU bus 1340. As noted, the software implements a routing protocol such as that mentioned above. Processor 1310 communicates with other systems of router 100 using an Ethernet communications mechanism via a 100Mbps Ethernet transceiver 1350. Ethernet transceiver 1350 is depicted in Fig. 13 as including a 100Mbps MAC 1351, a PHY/transceiver 1352, a transformer 1353 and a switch 1354. Switch 1354 provides a redundant connection to the other systems of router 100 to allow uninterrupted operation in the event of a communications failure.

System Switch

Fig. 14 illustrates an example of a system switch depicted as a system switch 1400, which can use an Ethernet-based communications, for example. In an Ethernet configuration, system switch 1400 manages the Ethernet connections from all level-1, level-2, route, and optional Wide Area Network (WAN) processors (not shown). System switch 1400 implements a high-speed, low-latency Ethernet switch that isolates local traffic to individual segments. The core of system switch 1400 is a switch matrix 1410. In one embodiment, switch matrix

-22-

1410 is an eight port bus that interconnects switch port controllers 1420(1)-(N), one or more high-speed interfaces (exemplified by a gigabit Ethernet switch port controller 1430), and expansion ports 1440(1)-(N). Each one of expansion ports 1440(1)-(N) communicates with a corresponding one of expansion buses 1450(1)-(N), respectively. Switch matrix 1410 is controlled by a processor 1460. Each copy of system Ethernet switch 1400 thus supports communications with level-1 processors, route processors, each I/O bay, and each matrix shelf processor. In Ethernet-based systems, these connections may be by 100Mbps or 10Mbps connections.

Software Architecture

In one embodiment, router 100 implements many functions in software to provide flexibility, support for communications protocols, and ease of implementation. The software architecture presented here forms a distributed management, control, and routing layer capable of spanning hundreds or thousands of nodes. The software architecture covers all protocol layers, management and control applications, and inter-node communication protocols and APIs.

The software modules described herein may be received by the various hardware modules of router 100, for example, from one or more computer readable media. The computer readable media may be permanently, removably or remotely coupled to the given hardware module. The computer readable media may non-exclusively include, for example, any number of the following: magnetic storage media including disk and tape storage media; optical storage media such as compact disk media (e.g., CD-ROM, CD-R, etc.) and digital video disk storage media; nonvolatile memory storage memory including semiconductor-based memory units such as FLASH memory, EEPROM, EPROM, ROM or application specific integrated circuits; volatile storage media including registers, buffers or caches, main memory, RAM, etc.; and data transmission media including computer network, point-to-point telecommunication, and carrier wave transmission media. In a UNIX-based embodiment, the software modules may be embodied in a file which may be a device, a terminal, a local or remote file, a socket, a network connection, a signal, or other expedient of communication or state change. Other new and various types of computer-readable media may be used to store and/or transmit the software modules discussed herein.

Overall Architecture

The software running the various processors of router 100 normally includes three major components: operating system, inter-processor and inter-node communications, and management and control applications. An important aspect of any software architecture is its underlying inter-process communications (IPC) mechanism.

IPCs that provide for the isolation of tasks are preferable. Such IPCs use message passing as their preferred communication. Message passing allows for full, but isolated interaction among tasks. To the rest of the system, a task, no matter how complex, is reduced to a simple producer and consumer of messages. It provides a set of well defined services, each accessed through one or more messages. Though sometimes visible to other tasks, in one embodiment, none of a given task's variables and structures should be accessible outside its

-23-

context. Limiting task interactions to message passing and keeping runtime variables private to each task allows individual software components to evolve independently and in parallel.

In order to keep code generic (i.e., system-and processor-independent), the message-based IPC should also provide a consistent application programming interface (API) that doesn't rely on any system-specific features or attributes. The API should have the same syntax and behavior, regardless of the underlying operating system, processor, or message-passing mechanism used. With certain generating systems, for example, message queues are used to implement the IPC, while on other kernels, pipes might be more appropriate. Preferably, then, the API should provide the following services to the application code:

1. Send message;
2. Receive a message;
3. Check for available messages; and
4. Name lookup and registration.

The last service, name lookup and registration, makes it possible for communicating entities to reference one another using names rather than task ID's, which are system-dependent.

15 Resource Manager

A resource manager (RM) is the software module responsible for collecting information about available resources and monitoring their status during normal system operation. A resource is used generically in this document to refer to any manageable hardware element that performs one or more system functions. The RM builds its resource list from unsolicited information it receives from other modules in the system, and from periodic keep-alive messages it exchanges with those modules. The RM, for example, is the first system application notified of card failures, insertions, and removals.

In one embodiment of router 100, there are two RM versions in the system. The first, which runs on the level-1 processor, is responsible for managing system resources and, in some cases, network-wide resources. The other version, which runs on level-2 processors, is responsible for managing resources in a single shelf. This multi-level hierarchy creates a flexible and expandable system where lower-level resource managers are custom designed for the specific shelf controlled.

The RM maintains information about a given resource in a structure called the Resource Control Block (RCB). The RCB consists of two main sections: a generic section, which is the same for all resources regardless of type, and a resource-specific section that varies according to resource type. All resource managers maintain a hierarchical list of resource control blocks that represents resources under their control. The list is referred to herein as the resource list and reflects the resources' hierarchy and their interdependencies. This allows the RM to determine, relatively quickly, the effect a given resource's failure has on other members of the hierarchy.

-24-

The router 100 preferably runs one or more versions of the Unix operating system on the level-1 processor and the level-2 processors (in the I/O and matrix shelves). Level-2 processors preferably run a real-time version of the Unix operating system (OS). Other processors (e.g., level-3, route, group, and matrix-node processors) preferably run a single task that does not require the services of an operating system or kernel. While
5 Unix operating systems are described herein as being preferable, any one of a number of operating systems may be used.

System Controller

The system controller is responsible for overall system management and control. The system controller uses a variety of protocols to communicate with other nodes in the network, including the operating system (OS).
10 Some of the protocols satisfy specific requirements (e.g. in a SONET based system, the transfer of OAM&P message across the SONET/SDH communications channels DCC), while others implement features, or functions, that are not part of the physical protocol used. To facilitate these functions, every router in a network is assigned an ID that uniquely identifies it within the network. The ID can also serve as a priority metric that determines the node's level within the hierarchy. However, the network can be configured to allow the user to override this by
15 manually assigning priorities to network nodes. The system controller supports a number of tasks that perform management, control, and routing functions, including resource management, OS interfacing, various network protocol servers, and operations, control, and intermediate system services.

Matrix Shelf Processor

The matrix shelf processor is responsible for the overall operation of a single main matrix shelf. It
20 communicates with the system controller, the route processor, and the microcontroller on each of the switch nodes, to provide local control and management for the shelf, including matrix configuration, diagnostics, and error reporting. The software on the matrix shelf processor preferably runs under a real-time Unix operating system. The RM on the matrix shelf processor is responsible for managing the hardware resources in its shelf. Like other resource managers in the system, the level-2 manager on this module uses a combination of hardware
25 and software to discover and maintain a list of available shelf resources. A protocol may be implemented to support such messaging.

In one embodiment, fault isolation is implemented by a dedicated task that is responsible for locating failures within the shelf. In a SONET based implementation, the software running on the shelf processor, with help from the microcontroller on the switch node, to determine(s) the quality of any of the input signals.

30 I/O Shelf Processor

Line Card Processor

The line card terminates an input signal from one of the other nodes in the network. For example, in a SONET-based implementation, a single SONET/SDH OC-48 signal is terminated by a line card, although other signal levels (OC-192, OC-12, and so on) may be supported. In one embodiment, the software consists of two

-25-

threads, one that runs in the background and is responsible for non-time critical tasks. The other thread, which runs at the interrupt level, is responsible for all real-time aspects of the software, including limited overhead processing, alarm detection and forwarding, and fault detection and recovery. The line card processor maintains a copy of its firmware and startup code onboard.

- 5 While particular embodiments of the present invention have been shown and described, it will be obvious to those skilled in the art that, based upon the teachings herein, changes and modifications may be made without departing from this invention and its broader aspects and, therefore, the appended claims are to encompass within their scope all such changes and modifications as are within the true spirit and scope of this invention. Furthermore, it is to be understood that the invention is solely defined by the appended claims.

WE CLAIM:

- 1 1. A network comprising:
2 a signal router configured to receive information carried by a first signal and to transmit said
3 information on a second signal, wherein
4 said second signal is selected from a plurality of signals,
5 said first and said second signals differ in at least one physical characteristic, and
6 said signal router is further configured to select said second signal from said plurality of
7 signals based on routing information obtained by said signal router.
- 1 2. The network of claim 1, wherein
2 said signal router is one of a plurality of such signal routers, each one of said plurality of such
3 signal routers is coupled to at least one other of said plurality of such signal routers,
4 said routing information is used to create a circuit between a first one of said plurality of such
5 signal routers and a second one of said plurality of such signal routers, and
6 said information is routed over said circuit.
- 1 3. The network of claim 2, wherein said first signal and said plurality of signals are
2 optical signals.
- 1 4. A signal router comprising:
2 a switching matrix configured to:
3 accept a first carrier, said first carrier carrying information,
4 select a second carrier from a plurality of carriers based on routing information
5 obtained by said signal router, said first and said second carriers differing
6 from one another in at least one physical characteristic, and
7 output said information on said second carrier.
- 1 5. The signal router of claim 4, wherein said first carrier is light of a first wavelength and said
2 second carrier is light of a second wavelength.
- 1 6. A signal router comprising:
2 an optical receiver configured to receive light of a first wavelength, wherein said light of said
3 first wavelength carries information;
4 an optical transmitter configured to transmit light of a second wavelength; and
5 a switching matrix, coupled to said optical receiver and said optical transmitter, and configured
6 to:

7 select said optical transmitter from a plurality of such optical transmitters, based on
8 routing information obtained by said signal router, and
9 transfer said information from said optical receiver to said optical transmitter.

1 7. The signal router of claim 6, wherein said signal router further comprises a group matrix,
2 wherein
3 said group matrix is coupled between said optical receiver, said optical transmitter, and said
4 switching matrix,
5 each one of said plurality of receivers supports a plurality of input channels, and
6 said group matrix is configured to perform optical line switching in order to allow selection of
7 a working channel from said plurality of input channels.

1 8. The signal router of claim 6, wherein wherein said optical receiver and said optical
2 transmitter are combined into an optical transceiver.

1 9. The signal router of claim 6, wherein each of said transmitters transmits light of a different
2 wavelength.

1 10. The apparatus of claim 33, wherein said signal path further comprises an input/output
2 section, coupled to said switching matrix.

1 11. The apparatus of claim 10, wherein said input/output section comprises:
2 a group matrix, coupled to said switching matrix; and
3 an input/output group, coupled to said group matrix.

1 12. The apparatus of claim 11, wherein said input/output group comprises a plurality of
2 optical line cards.

1 13. The apparatus of claim 33, wherein said control path comprises:
2 a local area network;
3 a first processor, coupled to said local area network and configured to manage said signal path;
4 a second processor, coupled to said local area network, and configured to determine said
5 routing information and generate switch configuration information; and
6 a third processor, coupled to said local area network and said switching matrix, and configured
7 to control said switching matrix based on said switch configuration information.

1 14. The apparatus of claim 13, wherein said switching matrix comprises a plurality of matrix
2 elements.

1 15. The apparatus of claim 33, wherein said control path comprises:
2 a system switch, coupled to said switching matrix;
3 a level-1 processor, coupled to said system switch;
4 a route processor, coupled to said system switch; and
5 a line rack, coupled to said system switch.

1 16. The apparatus of claim 15, wherein said line rack comprises:
2 a shelf switch, coupled to said system switch; and
3 a plurality of groups, coupled to said shelf switch.

1 17. The apparatus of claim 16, wherein each one of said plurality of groups comprises:
2 a first group switch, coupled to said shelf switch;
3 a second group switch, coupled to said shelf switch; and
4 a plurality of line cards, coupled to said first and second group switches.

1 18. A method for routing information across a network, the method comprising:
2 obtaining routing information, wherein the network comprises a plurality of nodes, and each
3 one of said plurality of nodes is coupled to at least one other of said plurality of nodes
4 by one of a plurality of links;
5 receiving a first signal at a first one of said plurality of nodes, wherein said first signal carries
6 the information;
7 selecting a second signal from a plurality of signals based on said routing information, wherein
8 said first and second signals differ in at least one physical characteristic;
9 transferring the information from said first signal to said second signal; and
10 transmitting said second signal from said first one of said plurality of nodes.

1 19. The method of claim 18, further comprising:
2 creating a physical path between said plurality of nodes using the obtaining, receiving,
3 selecting, and transmitting step; and
4 wherein said transmitting step comprises routing said signal over a virtual path created over
5 said physical path.

1 20. The method of claim 19, wherein said virtual path is a telecommunications circuit.

1 21. The network of claim 2, wherein said routing information is obtained from at least
2 one other of said plurality of such signal routers using in-band communications.

1 22. The network of claim 2, wherein said routing information is obtained by sending a
2 message to at least one other of said plurality of such signal routers, said message causing said at least
3 one other of said plurality of such signal routers to respond by sending a response containing topology
4 information stored on said at least one other of said plurality of such signal routers, said topology
5 information representing a topology of at least a portion of said network.

1 23. The network of claim 2, wherein said signal router comprises:
2 a receiver configured to receive said first signal;
3 a transmitter configured to transmit said second signal; and
4 a switching matrix, coupled to said receiver and said transmitter, and configured to:
5 select said transmitter from a plurality of such transmitters, based on said routing
6 information, in order to perform said selection of said second signal from
7 said plurality of signals, and
8 transfer said information from said receiver to said transmitter.

1 24. The network of claim 23, wherein said physical characteristic is wavelength.

1 25. The network of claim 23, wherein said physical characteristic is phase.

1 26. The signal router of claim 4, wherein said signal router comprises:
2 a receiver configured to receive said first carrier;
3 a transmitter configured to transmit said second carrier; and
4 a switching matrix, coupled to said receiver and said transmitter, and configured to:
5 select said transmitter from a plurality of such transmitters, based on said routing
6 information, in order to perform said selection of said second carrier, and
7 transfer said information from said receiver to said transmitter.

1 27. The signal router of claim 26, wherein said first carrier is a first optical signal and
2 said second carrier is a second optical signal.

1 28. The signal router of claim 27, wherein said physical characteristic is wavelength.

1 29. The signal router of claim 4, wherein said routing information is obtained from at
2 least one other of a plurality of such signal routers using in-band communications.

1 30. The signal router of claim 4, wherein said routing information is obtained by sending
2 a message to at least one other of a plurality of such signal routers, said message causing said at least
3 one other of said plurality of such signal routers to respond by sending a response containing topology
4 information stored on said at least one other of said plurality of such signal routers, said topology
5 information representing a topology of at least a portion of said network.

1 31. The signal router of claim 6, wherein said routing information is obtained from at
2 least one other of a plurality of such signal routers using in-band communications.

1 32. The signal router of claim 6, wherein said routing information is obtained by sending
2 a message to at least one other of a plurality of such signal routers, said message causing said at least
3 one other of said plurality of such signal routers to respond by sending a response containing topology
4 information stored on said at least one other of said plurality of such signal routers, said topology
5 information representing a topology of at least a portion of said network.

1 33. A signal router comprising:
2 a control path; and
3 a signal path, coupled to said control path, wherein said signal path comprises a switching
4 matrix, wherein said switching matrix comprises at least one stage and configured to:
5 accept a first carrier, said first carrier carrying information,
6 select a second carrier from a plurality of carriers based on routing information from
7 said control path, said first and said second carriers differing from one
8 another in at least one physical characteristic, and
9 output said information on said second carrier.

1 34. The network of claim 18, wherein said first signal and said plurality of signals are
2 optical signals and said physical characteristic is wavelength.

1 35. The network of claim 18, wherein said obtaining comprises:
2 requesting said routing information from at least one other of said plurality of nodes; and
3 receiving said routing information from said at least one other of said plurality of nodes,
4 wherein said requesting and said receiving employ in-band communications

1 36. The network of claim 18, said obtaining comprising:
2 sending a message from said first one of said plurality of nodes to at least one other of said
3 plurality of nodes,

4 sending a response from said at least one other of said plurality of nodes to said first one of
5 said plurality of nodes in response to said message, wherein said response contains
6 topology information stored on said at least one other of said plurality of nodes, and
7 said topology information represents a topology of at least a portion of said network.

1 37. The network of claim 18, wherein said first signal and said plurality of signals are
2 optical signals.

1 38. A signal router comprising:
2 optical receiving means for receiving light of a first wavelength, wherein said light of said first
3 wavelength carries information;
4 optical transmitting means for transmitting light of a second wavelength; and
5 switching means, coupled to said optical receiving means and said optical transmitting means,
6 for:
7 selecting said optical transmitting means from a plurality of such optical transmitting
8 means, based on routing information obtained by said signal router, and
9 transferring said information from said optical receiving means to said optical transmitting
10 means.

1 39. The signal router of claim 38, wherein said signal router further comprises:
2 optical line switching means for selecting a working channel from a plurality of input channels,
3 wherein
4 said receiving means receives said first plurality of input channels,
5 said light of said first wavelength is one of said plurality of input channels,
6 said optical line switching means is coupled between said optical receiving means and
7 said switching means; and
8 optical splitting means for creating a plurality of output channels, wherein
9 said optical splitting means is coupled between said switching means and said optical
10 transmitting means.

1 40. A signal router comprising:
2 signal path means for routing a plurality of input signals to a plurality of output signals, said
3 signal path means comprising:
4 receiving means for receiving one of said plurality of input signals, said one of said
5 plurality of input signals carrying information,
6 switching means for:
7 accepting said information from said receiving means, said receiving means
8 coupled to one of a plurality of inputs of said switching means,

9 selecting one of a plurality of outputs of said switching means based on
10 routing information, and
11 outputting said information at said one of said plurality of outputs, and
12 transmitting means for transmitting said information on one of said plurality of output
13 signals, said one of said plurality of outputs coupled to said transmitting
14 means; and
15 control path means for controlling said signal path means, said control path means
16 obtaining said routing information.




FIG. 1A

2/14




FIG. 1B

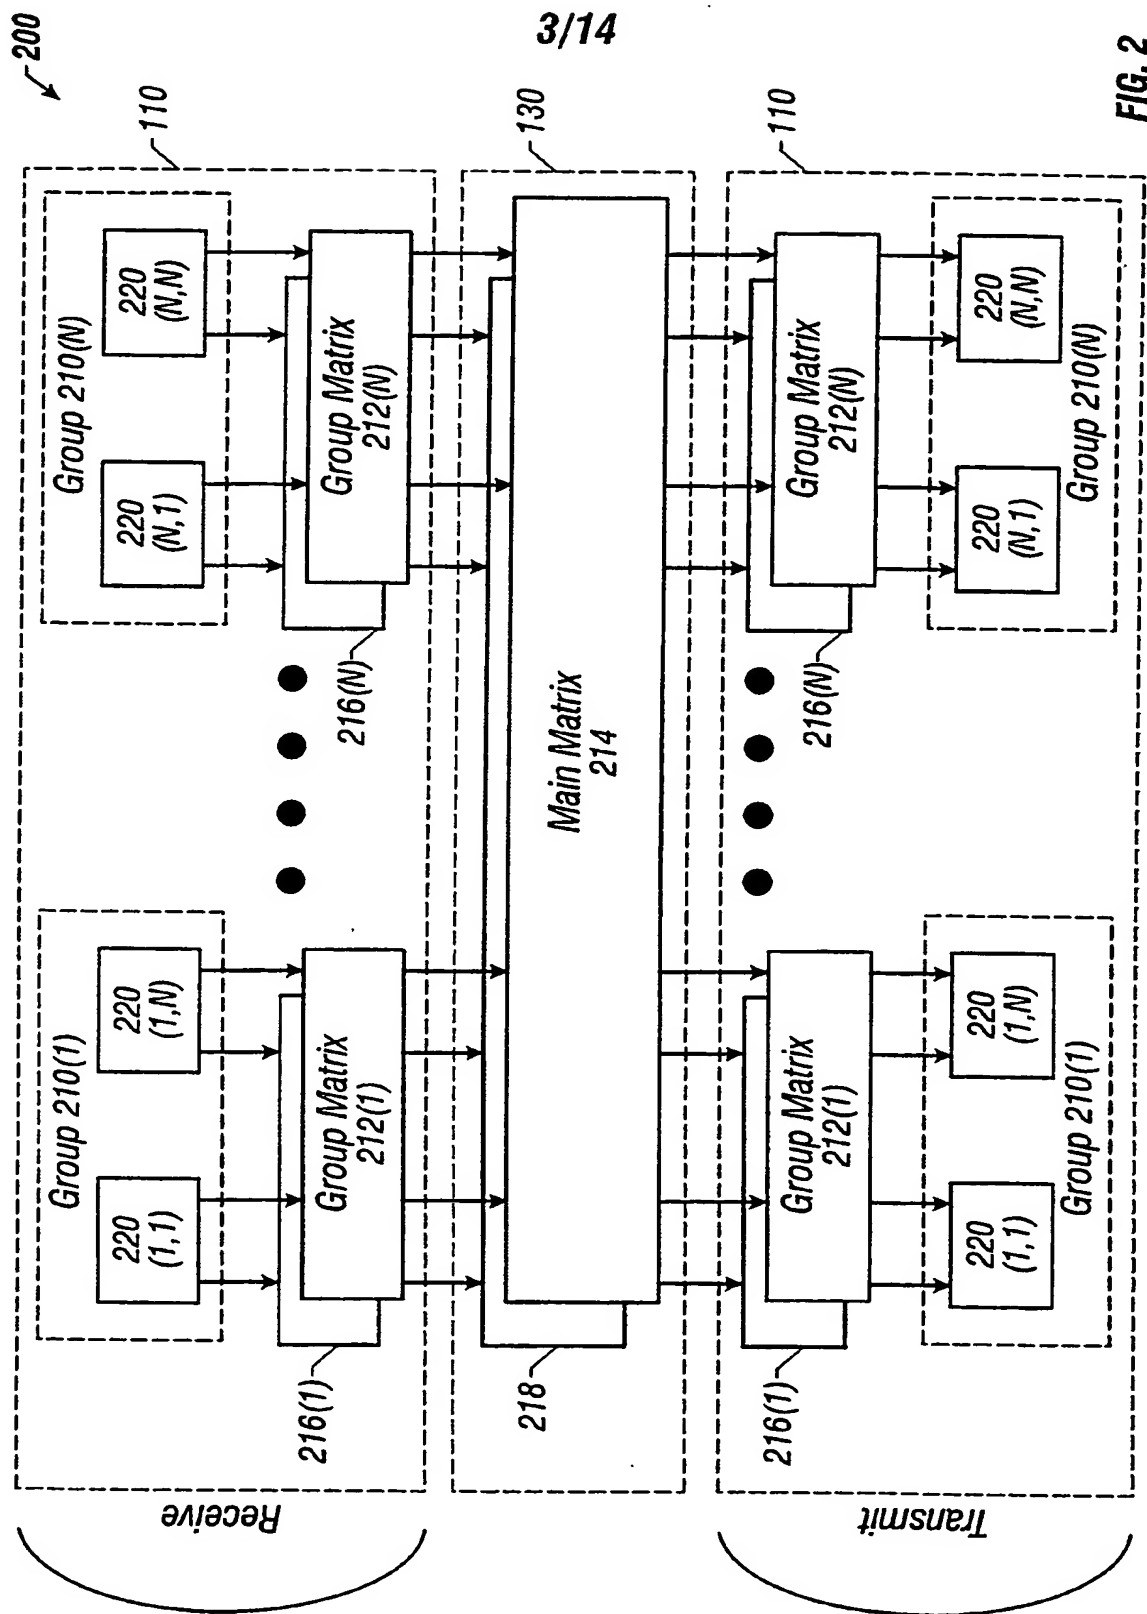



FIG. 2

4/14

FIG. 3






FIG. 4




FIG. 5

7/14

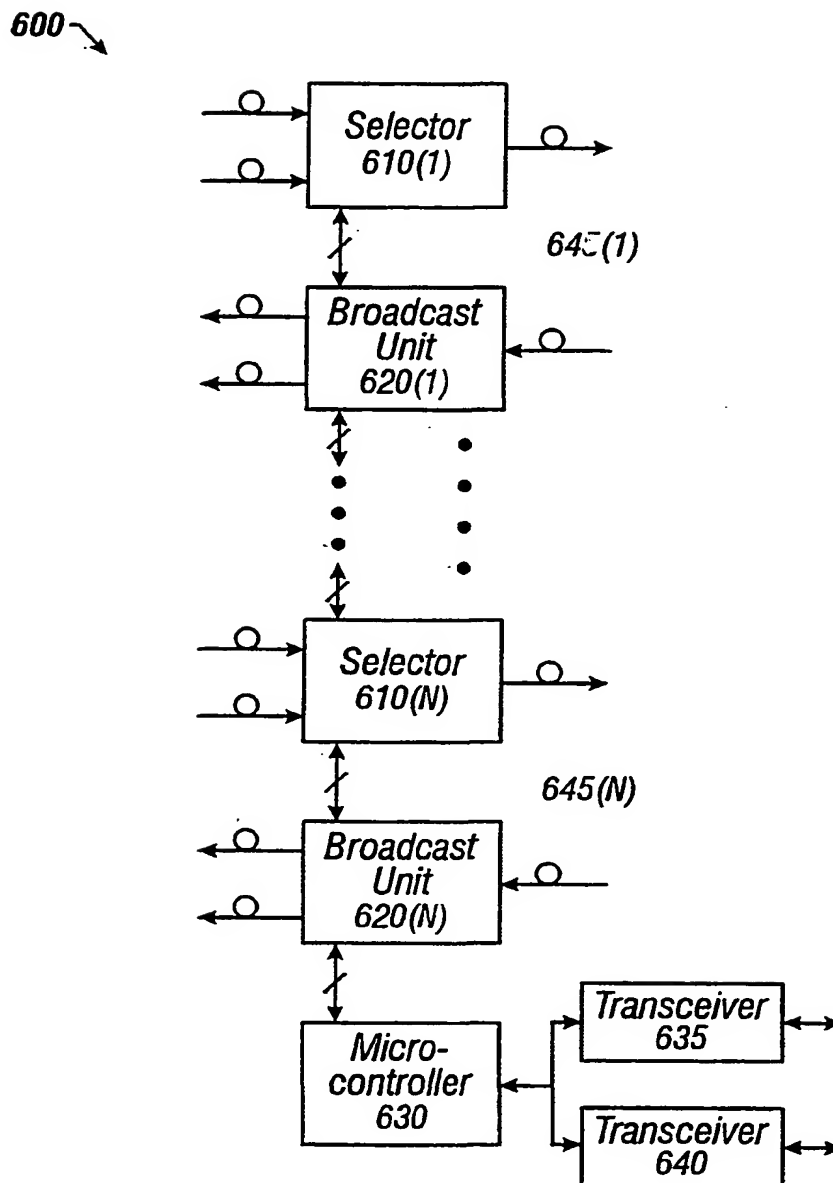


FIG. 6

8/14

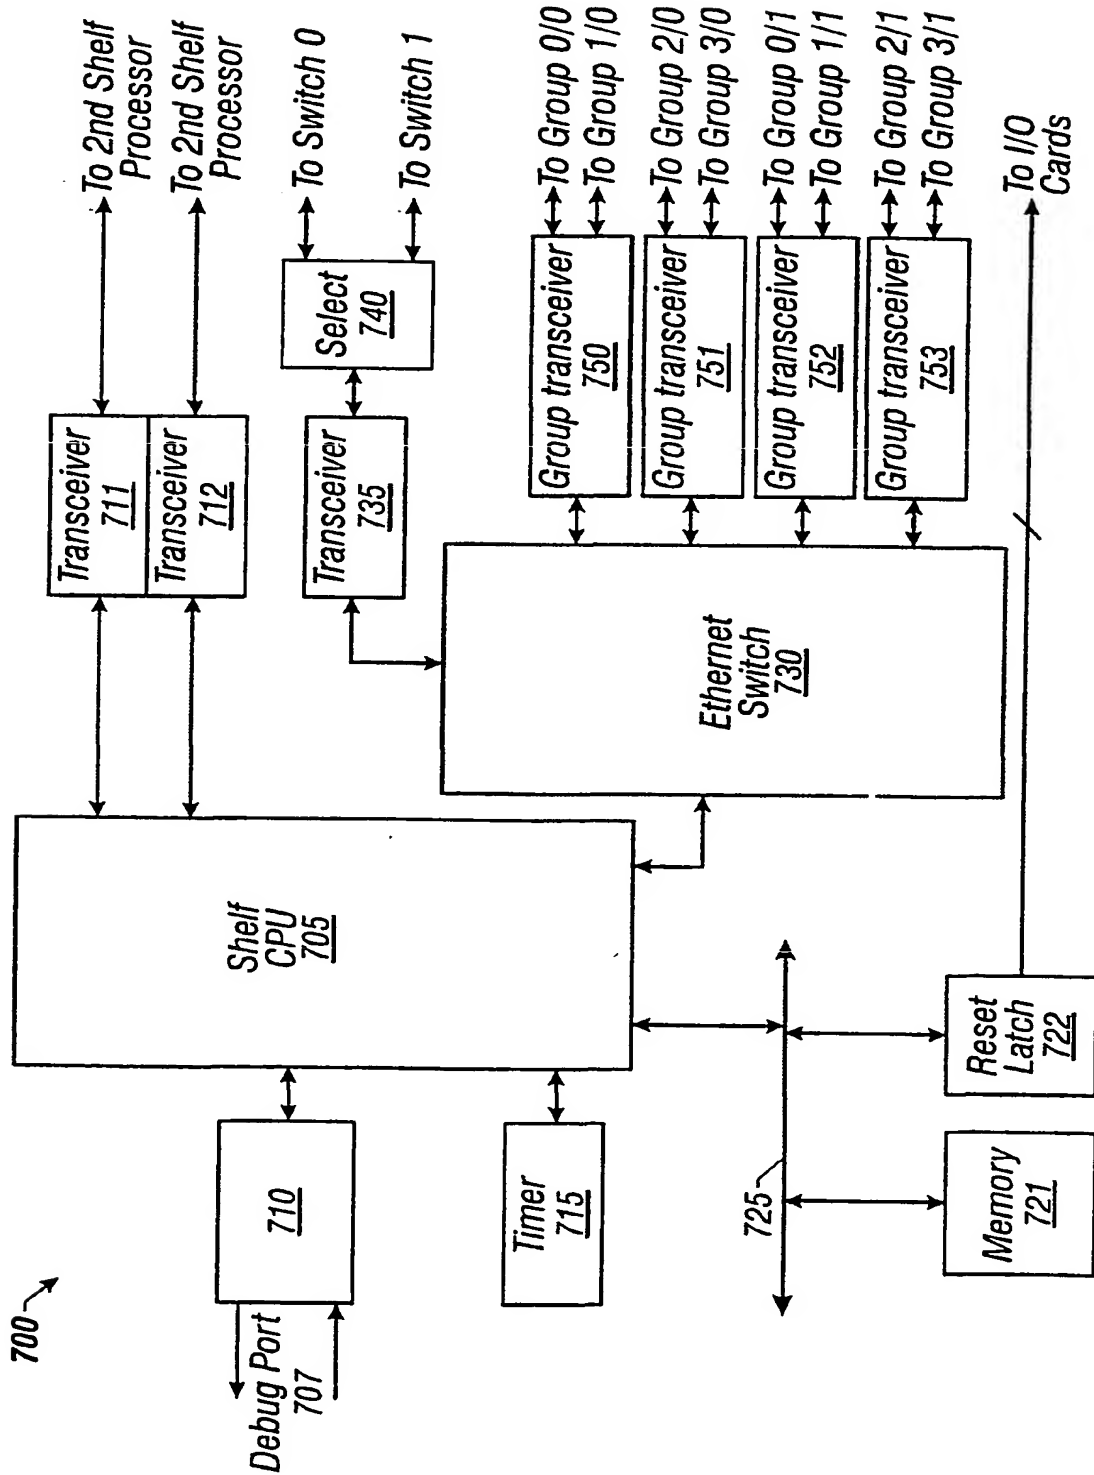


FIG. 7

9/14

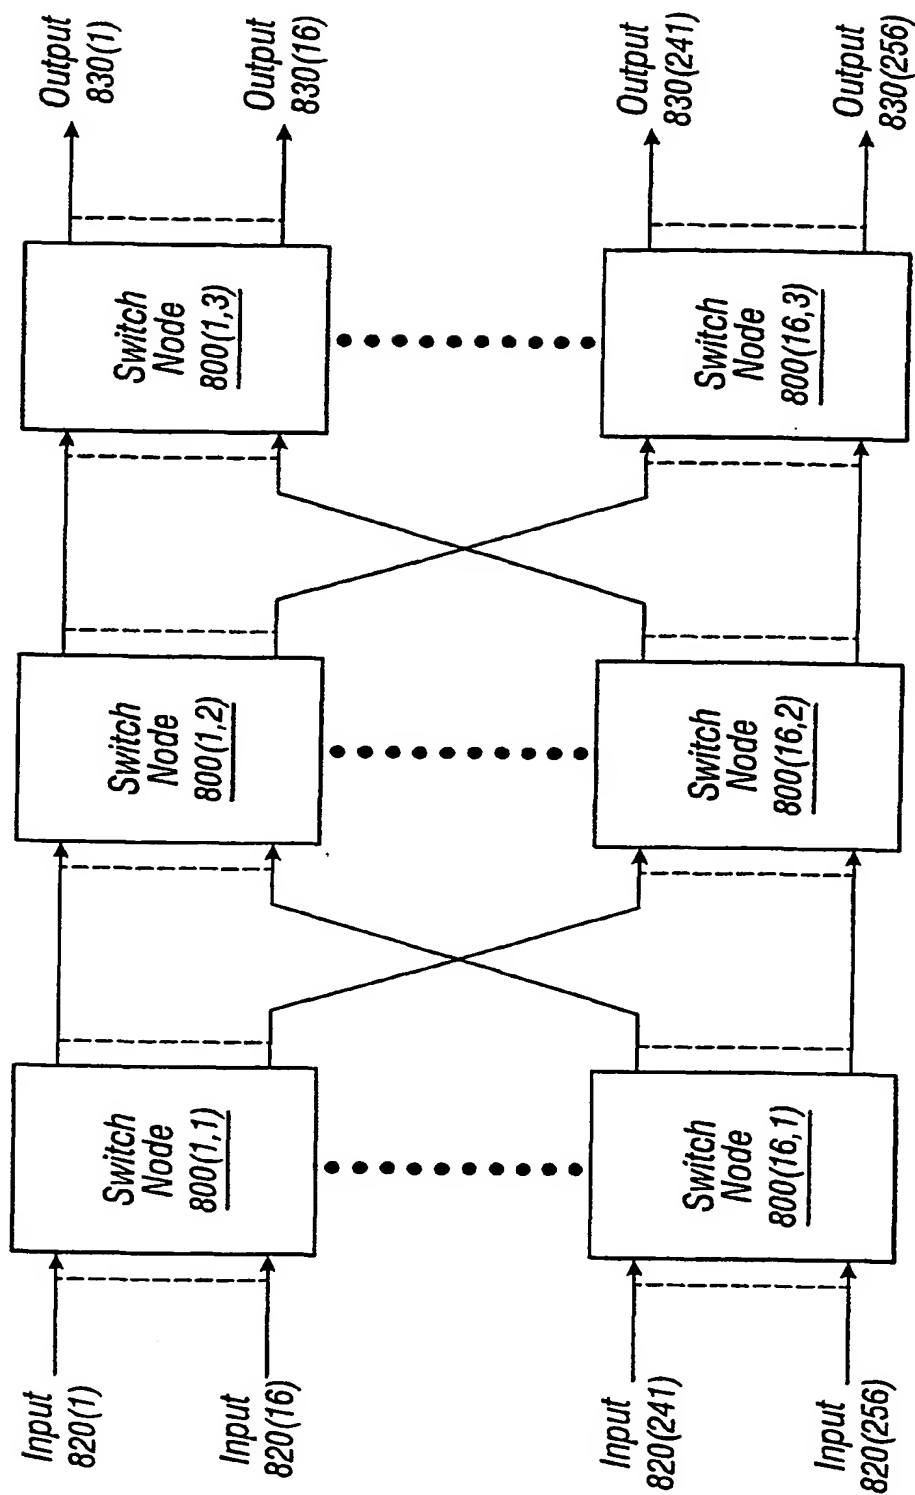


FIG. 8

10/14

Matrix Rack 900 ↘

800(1,3)		800(1,2)		800(1,1)	
800(2,3)		800(2,2)		800(2,1)	
800(3,3)		800(3,2)		800(3,1)	
800(4,3)		800(4,2)		800(4,1)	
800(5,3)		800(5,2)		800(5,1)	
800(6,3)		800(6,2)		800(6,1)	
800(7,3)		800(7,2)		800(7,1)	
800(8,3)		800(8,2)		800(8,1)	
910(6)	910(5)	910(4)	910(3)	910(2)	910(1)
800(9,3)		800(9,2)		800(9,1)	
800(10,3)		800(10,2)		800(10,1)	
800(11,3)		800(11,2)		800(11,1)	
800(12,3)		800(12,2)		800(12,1)	
800(13,3)		800(13,2)		800(13,1)	
800(14,3)		800(14,2)		800(14,1)	
800(15,3)		800(15,2)		800(15,1)	
800(16,3)		800(16,2)		800(16,1)	

FIG. 9

11/14




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/00970

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04Q11/00 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 802 697 A (FUJITSU LTD) 22 October 1997 (1997-10-22)	1,4,6, 18,33, 38,40
A	column 8, line 4 -column 9, line 31	2,3,5, 7-17, 19-32, 34-37,39
	column 43, line 30 -column 46, line 10	
A	US 5 777 761 A (FEE JOHN A) 7 July 1998 (1998-07-07) column 3, line 38 -column 4, line 19 column 4, line 60 -column 5, line 6 column 5, line 54 -column 6, line 2	1-40
	-/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "8" document member of the same patent family

Date of the actual completion of the international search

31 May 2000

Date of mailing of the international search report

20/06/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Meurisse, W

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/00970

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 97 24901 A (MCI COMMUNICATIONS CORP) 10 July 1997 (1997-07-10) page 6, line 8 -page 7, line 15 page 9, line 5 - line 20 ---	21,22
A	EP 0 752 795 A (IBM) 8 January 1997 (1997-01-08) column 8, line 22 -column 9, line 17 ---	21,22 -
A	US 5 657 449 A (OSAKI YOSHIRO) 12 August 1997 (1997-08-12) column 5, line 38 -column 7, line 8 column 8, line 17 - line 37 column 9, line 16 - line 33 ---	13,15
A	WO 97 04397 A (ASCOM NEXION INC) 6 February 1997 (1997-02-06) page 7, line 26 -page 8, line 25 page 9, line 21 - line 31 page 12, line 12 - line 23 page 14, line 9 - line 23 ---	13,15
A	KOMINE,CHUJO,OGURA,MIYAZAKI,SOEJIMA: "A distributed restoration algorithm for multiple-link and node failures of transport networks" IEEE, GLOBECOM '90, April 1990 (1990-04), pages 459-463, XP002139230 ---	1-40
A	ZHANG,WEI,QIAO: "On fundamental issues in IP over WDM Multicast" IEEE, COMPUTER COMMUNICATIONS AND NETWORKS, September 1999 (1999-09), pages 84-90, XP002139231 -----	1-40

INTERNATIONAL SEARCH REPORT

Information on patent family members

Internal Application No.

PCT/US 00/00970

Patent document cited in search report		Publication date	Patent family member(s)	Publication dat
EP 0802697	A	22-10-1997	JP 10004418 A	06-01-1998
			CN 1171683 A	28-01-1998
US 5777761	A	07-07-1998	US 5731887 A	24-03-1998
			CA 2241106 A	17-07-1997
			EP 0870412 A	14-10-1998
			WO 9725829 A	17-07-1997
WO 9724901	A	10-07-1997	US 5914798 A	22-06-1999
			CA 2241895 A	10-07-1997
			EP 0870413 A	14-10-1998
EP 0752795	A	08-01-1997	US 5781537 A	14-07-1998
			JP 9036892 A	07-02-1997
US 5657449	A	12-08-1997	JP 4100497 A	02-04-1992
			CA 2049476 A, C	21-02-1992
WO 9704397	A	06-02-1997	AU 6500796 A	18-02-1997
			AU 6500896 A	18-02-1997
			AU 6500996 A	18-02-1997
			AU 6501096 A	18-02-1997
			AU 6501496 A	18-02-1997
			AU 6501696 A	18-02-1997
			AU 6501796 A	18-02-1997
			AU 6501996 A	18-02-1997
			AU 6502096 A	18-02-1997
			AU 6502496 A	18-02-1997
			AU 6502596 A	18-02-1997
			AU 6502696 A	18-02-1997
			AU 6502796 A	18-02-1997
			AU 6503196 A	18-02-1997
			AU 6503296 A	18-02-1997
			AU 6503396 A	18-02-1997
			AU 6503496 A	18-02-1997
			AU 6503596 A	18-02-1997
			AU 6503696 A	18-02-1997
			AU 6503796 A	18-02-1997
			AU 6549196 A	18-02-1997
			AU 6549296 A	18-02-1997
			AU 6648496 A	18-02-1997
			AU 6648796 A	18-02-1997
			AU 6712496 A	18-02-1997
			AU 6712596 A	18-02-1997
			AU 6761896 A	18-02-1997
			AU 6762096 A	18-02-1997
			EP 0845181 A	03-06-1998
			EP 0839420 A	06-05-1998
			EP 0839419 A	06-05-1998
			EP 0872086 A	21-10-1998
			EP 0839421 A	06-05-1998
			EP 0839422 A	06-05-1998
			JP 11510323 T	07-09-1999
			JP 2000501897 T	15-02-2000
			JP 11510324 T	07-09-1999
			JP 2000501900 T	15-02-2000
			JP 2000501901 T	15-02-2000
			JP 11510003 T	31-08-1999

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/00970

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9704397 A		JP 2000501902 T	15-02-2000
		JP 11510004 T	31-08-1999
		JP 11510005 T	31-08-1999
		JP 11510006 T	31-08-1999
		JP 11511303 T	28-09-1999
		JP 11510007 T	31-08-1999
		JP 11510008 T	31-08-1999
		JP 11510009 T	31-08-1999
		JP 11510010 T	31-08-1999
		JP 11510011 T	31-08-1999